



Population coding under normalization

Dario L. Ringach

Departments of Neurobiology and Psychology, Jules Stein Eye Institute, David Geffen School of Medicine, University of California, Los Angeles, CA 90095, USA

ARTICLE INFO

Article history:

Received 13 October 2009

Received in revised form 9 December 2009

Keywords:

Normalization
Gain control
Population coding
Optimal coding
Visual cortex
Electrostatic potential
Thomson problem
Minimum energy

ABSTRACT

A common computation in visual cortex is the divisive normalization of responses by a pooled signal of the activity of cells within its neighborhood. From a geometrical point of view, normalization constrains the population response to high-contrast stimuli to lie on the surface of a high-dimensional sphere. Here we study the implications this constraint imposes on the representation of a circular variable, such as the orientation of a visual stimulus. New results are derived for the infinite dimensional case of a homogeneous populations of neurons with identical tuning curves but different orientation preferences. An important finding is that the ability of the population to discriminate between any two orientations depends exclusively on the Fourier amplitude spectrum of the orientation tuning curve. We also study the problem of encoding by a finite set of neurons. A central result is that, under normalization, optimal encoding can be achieved by a finite number of neurons with heterogeneous tuning curves. In other words, increasing the number of neurons in the population does not always allow for an improved population code. These results are used to estimate the number of neurons involved in the coding of orientation at one position in the visual field. If the cortex were to code orientation optimally, we find that a small number (~ 4) of neurons should suffice.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

The gain control model extended the classic linear–nonlinear model of simple cells (Movshon, Thompson, & Tolhurst, 1978) to account for a number of response properties, including response saturation, phase advance of temporal responses with contrast, and the results of masking experiments using plaid stimuli (Bonds, 1989; Carandini, Heeger, & Movshon, 1997; Heeger, 1992; Robson, 1988). The basic idea of the model is that an initial set of responses, provided by direct input from other brain regions, gets normalized (divided) by a pooled signal of neuronal activity in the neighborhood a cell (Fig. 1).

Progress over the last decade has shown that such normalization is a widespread computation in the brain; it can be found in the retina (Benardete, Kaplan, & Knight, 1992; Shapley & Victor, 1979a, 1979b, 1981; Solomon, Lee, & Sun, 2006), the lateral geniculate nucleus (Bonin, Mante, & Carandini, 2005, 2006), primary visual cortex (Carandini et al., 1997; Heeger, 1992; Ringach & Malone, 2007; Rust, Schwartz, Movshon, & Simoncelli, 2005), area MT (Simoncelli & Heeger, 1998), and area IT (Zoccolan, Cox, & DiCarlo, 2005). Furthermore, normalization models appear to account well for the modulatory effects of attention (Reynolds, Chelazzi, & Desimone, 1999; Reynolds & Heeger, 2009).

The prevalence of normalization in the nervous system must surely reflect the fact that it evolved to address a problem that arises

at different stages of processing (Douglas & Martin, 2004, 2007). From a theoretical point of view, this observation prompts a number of interesting questions. What basic principles of signal processing would lead to normalization of responses to be a critical component of neural computation? How does normalization impact the way stimuli can be encoded and processed? What are the computational capabilities of networks of normalized populations?

Some recent studies have considered how normalization modifies the statistical dependencies of neural activity, and have put forward the idea that normalization may serve to optimize the representation of natural signals (Fairhall, Lewen, Bialek, & de Ruyter Van Steveninck, 2001; Olshausen & Field, 1996a, 1996b, 2004; Ruderman & Bialek, 1994; Schwartz & Simoncelli, 2001). Others have noted that normalization may also serve a role in decoding the activity of neuronal populations (Deneve, Latham, & Pouget, 1999). Here we take complementary approach and ask not why, but *how* is that the representation of information constrained when carried by the signals of a normalized pool of neurons? Under what conditions are such representations optimal?

We study this problem in two scenarios that are simple enough to allow the derivation of theoretical results. First, we consider the encoding performed by a homogeneous set of neurons with identical tuning curves differing only in their preferred orientation. The main object of study is the *information tuning curve* (Kang, Shapley, & Sompolinsky, 2004) which specifies the ability of the population to discriminate between any two orientations. When the number of neurons tends to infinity, closed-form calculations can be

E-mail address: dario@ucla.edu




Fig. 1. Normalization of population responses. Front end linear receptive fields are followed by a half-rectifier (Movshon et al., 1978), and the resulting responses are normalized by the term $\sqrt{\sigma^2 + \|r\|^2}$. (Figure modified from Carandini et al. (1997).)

performed that clarify the constraints imposed by normalization onto the information tuning curve. In particular, it is shown that the information tuning curve is determined by the Fourier *amplitude* spectrum of the tuning curve. Surprisingly, the result can be used to show that the average discrimination performance of the population (measured as the average $(d')^2$ across all possible orientation pairs) depends exclusively on the *mean* of the tuning curve of neurons and not at all on its *shape*. This has important consequences for experiments in perceptual learning that attempt to gauge how neuronal populations change during the learning process. Namely, geometric properties of local shape of the tuning curve (such as its bandwidth at half-height or maximum slope) are not the best choice when evaluating how a population's ability to discriminate between any two orientations. Instead, its Fourier amplitude spectrum is better suited, as it encodes all the information in the information tuning curve.

Second, we study the more complex situation of finite dimensional cases, where neurons are also allowed to have different tuning curves. We ask how can we best represent a circular variable in an optimal fashion in this case. The main finding is, at first, counterintuitive: normalization causes the optimal encoding to be attainable using a *finite* number of neurons. In other words, increasing the number of neurons in a normalized population does not always allow for an improved population code.

We then show how these results can be used to estimate the number of neurons involved in the coding of orientation at any one position in the visual field from experimental data. We find that if the cortex were to code orientation optimally, a small number (~ 5) of normalized neurons would suffice.

These findings demonstrate that normalization imposes important constraints on the coding of information and, giving its incidence in cortical circuits, it should be incorporated as an integral component in formal models of population coding.

2. Results

2.1. A geometric view of normalization

The normalization model we adopt is one that has been used widely to model the responses of simple cells in primary visual cortex (Carandini et al., 1997; Heeger, 1992):

$$\hat{r}_i = r_i / \sqrt{\sigma^2 + \|r\|^2} \quad i = 0, 1, \dots, N-1 \quad (1)$$

Here r_i represents the initial response of neuron i to a stimulus, $r = (r_0, r_1, \dots, r_{N-1})$ is a vector describing the response of the N cells in the population, $\|r\|$ is the Euclidian norm of the vector, σ is the semi-saturation constant, and $\hat{r} = (\hat{r}_0, \hat{r}_1, \dots, \hat{r}_{N-1})$ is the vector of normalized responses. Our first observation is simple: for strong stimuli, for which $\|r\| \gg \sigma$, the normalized population response lies on a hyper-sphere, S^{N-1} . This means that in general the problem of encoding under normalization becomes one of defining a map from a given stimulus space to a high-dimensional sphere, the dimension of which is determined by the number of neurons at hand. A few examples demonstrate how this fact leads to some interesting theoretical problems.

Consider first the coding of a circular variable, which we will discuss in detail below. This is a problem confronted by the nervous system in several contexts, including the coding of the orientation or direction of motion of a visual stimulus, the color hue of a surface patch, wind direction, reaching direction, and heading direction with respect to magnetic north. In all these cases, the domain of the stimulus can be associated to unit circle. Any point on the circle defines a value of the variable under consideration (Fig. 2a). Thus, the problem of encoding a circular variable by a normalized population of N neurons amounts to defining a map that takes the unit circle into the a $(N-1)$ -sphere,

$$f : S^1 \rightarrow S^{N-1}. \quad (2)$$

Studying the property of these maps, and how one could obtain some that are optimal in some sense, is the central topic of our study.

Another interesting situation arises when we consider the joint coding of the orientation and the spatial phase of a sinusoidal grating (Fig. 2b). This stimulus space represents a Klein bottle, a fact that may not be entirely evident at first (Carlsson, Ishkhanov, DaSilva, & Zomordian, 2008; Singh et al., 2008; Swindale, 1996; Tanaka, 1995). One way to show this is by plotting the individual gratings corresponding to the various combinations of the parameters (Fig. 2b). One can then see that gratings on the top and bottom rows are identical to each other, as indicated by the matching directions of the two red arrows. If we wanted to match these two arrows one on top of the other all we need to do is roll the rectangle of gratings into a cylinder. On the other hand, the stimuli on the left and right columns are the same but reversed in order, as indicated by the directions of the blue arrows. To match the directions of the blue arrows once we have rolled the space into a cylinder we would need to puncture the cylinder in 3D space, but it turns out that it can be done in 4D without puncturing (Fig. 2b). The resulting object is a Klein bottle. Thus, the problem of mapping the orientation and spatial phase of a grating onto the normalized population of cells amounts to defining an embedding of the Klein bottle into the $(N-1)$ -sphere.

A final example is the joint coding of orientation and color hue (Johnson, Hawken, & Shapley, 2008) (Fig. 2c). Here, as one can infer by the direction of the arrows, the resulting object is a torus, which is obtained by rolling the space into a cylinder and then gluing the ends together. Implementing a population code that maps orientation and color hue to a normalized population of neurons is thus equivalent to the embedding of a torus in a high-dimensional sphere.

These examples illustrate that interesting geometrical and topological problems arise when we consider the representation of visual information via normalized population of neurons. In what follows, we take a first step to gain some insight into the constraints imposed by normalization by considering a very simple case of orientation tuning (Benyishai, Baror, & Sompolinsky, 1995; Salinas & Abbott, 1994; Seung & Sompolinsky, 1993).

Download English Version:

<https://daneshyari.com/en/article/6203834>

Download Persian Version:

<https://daneshyari.com/article/6203834>

[Daneshyari.com](https://daneshyari.com)